


CONNEXIONS



The Interoperability Report

May 1987

Volume 1, No. 1

ConneXions - The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

Broadcast and Multicast.....	2
Conference report.....	4
Heard the GOSIP?.....	5
ISODE revealed.....	8
<i>gated</i> - A user's review.....	13
NetBIOS RFCs issued.....	15
TCP at ACM SIGCOMM....	15

ConneXions is published by Advanced Computing Environments, 21370 Vai Avenue, Cupertino, CA 95014, USA 408-996-2042.

© 1987 Advanced Computing Environments. Quotation with attribution is encouraged.

From the Editor

This is the first **regular** issue of *ConneXions - The Interoperability Report*. Our Premiere Issue promised articles on various aspects of interoperability. In this issue we explore some of the politics and opinions that affect our industry. The ISO versus TCP argument has tended to be an "us against them" battle with few signs of "harmonization" or "peaceful migration". Marshall Rose of Northrop Research and Technology Center in Palos Verdes, California has been working on the ISO Development Environment (ISODE). We present some of his thoughts on co-existence and migration in this issue. We also find out about GOSIP and the controversy surrounding it.

Broadcasting and multicasting in an internet raises several complex issues. We asked Vint Cerf of the Corporation for National Research Initiatives to introduce this topic.

The NSFnet project is an ambitious undertaking requiring interconnection of several kinds of networks. In this environment it is crucial to maintain proper routing management between all the gateways. The *gated* program, as reviewed in this issue by Sergio Heker from the Jon von Neumann National Supercomputer Center, provides a mechanism for managing gateways which speak a variety of interior gateway protocols.

The TCP/IP Interoperability Conference which was held in Monterey March 16 - 19, attracted nearly 700 people from the research, vendor, and user communities. We received a great deal of positive feedback from the attendees, plus some useful critique. In general, the critique fell into two broad categories: "This is too technical!" and "This isn't technical enough!" Clearly, this illustrates the diversity of our audience. Once again, thank you for your comments and suggestions. With them we can plan future events which will fit the needs of everyone working in the field of interoperability.

Your input for *ConneXions* is also encouraged. Our newsletter should be regarded as a forum for the community by the community. Keep in touch!

Broadcasting and Multicasting

by Vint Cerf, Corporation for National
Research Initiatives

There is a substantial current interest in the topics of broadcasting and multicasting in local area and larger networking environments. To some extent, this interest has been sparked by the recognition of the usefulness of these techniques in applications which have been largely confined to local area nets until recently. This essay explores some of the issues arising as these facilities are extended beyond the local area network.

What is it and why do it?

At its simplest, broadcasting is a technique for transmitting information to an unknown recipient. Many of the applications are of the form: *"Is there a doctor in the house?"* as in the request for a "breath of life" packet when a diskless workstation is first powered up. A "bootserver" on the network can listen for such broadcasts and provide a packet sequence in return which is used to initialize the operation of the workstation.

Other applications involve servers such as time servers which respond to the question *"What time is it?"* or gateway servers which respond to queries such as *"Is there a gateway to other networks available?"*

This style of operation is based on the notion that work stations should be able to start in a kind of "Garden of Eden" state (i.e. no knowledge at all) and become functional through access to the local area network (which, I suppose, plays the role of snake in this little allegory...).

Too much of a good thing

A problem with this particular line of reasoning is that one can end up with too much of a good thing. If every protocol makes use of broadcast in any quantity, there will be a substantial amount of traffic for each workstation to examine since each workstation has to examine every broadcast packet. Within each vendor offering, broadcasting appears to have been used fairly judiciously, but users are beginning to discover that bringing up multi-vendor systems on a common local area net generates broadcasts for every workstation to examine which were not intended or anticipated by any particular vendor.

Multicast is broadcast to a few

Multicasting is a restricted form of broadcasting. In this case, a packet is marked as destined for a particular multicast address and only those workstations which have joined that particular multicast group need to examine it. Of course, at the hardware level, every packet must be examined for relevance to the attached workstation, but a table of specific and multicast addresses, set by the workstation, can determine which packets pass the fast hardware filter.

The advantage of multicasting is that it potentially interrupts fewer workstations in the course of their duties. However, the workstation(s) must know which multicast groups to join, so this will require either some advance, out-of-band exchanges or perhaps some exchanges with one or more servers which are discovered by means of broadcasting.

Bindings For many computer communications applications, an essential element for success is for two or more correspondents to discover each other and to communicate. This is not always a trivial matter. First, there is the business of knowing how to refer to your correspondent(s). John Shoch's now famous essay on Names, Addresses and Routes, makes this point most clearly. *

Names must be transformed into addresses and addresses into routes before one can communicate. Broadcasting can eliminate the need for the routing step and even the addressing step, but at a cost which may be so prohibitive that it can be afforded only for a very few, crucial applications, or in a limited networking environment.

Indeed, where it is difficult or impossible to discover the translation of a name into an address, designers often resort to broadcast as an alternative. The NetBIOS concepts, which are based on names that are bound to addresses through a process called discovery, are firmly rooted in the idea of broadcasting.

Name Servers Where it is prohibitive to use broadcast methods to bind names to addresses, one must seek alternatives. In the Internet, the concept of Name Server represents one such alternative. While it might be necessary to rely on a (local?) broadcast to discover the address of one or more Name Servers, after that, it is hoped that at least one Name Server will continue to be available and can provide needed information (rather like Dear Abby and Dear Ann Landers, networks need redundant advice-givers).

Internet broadcasts can be dangerous Broadcasting throughout the length and breadth of the Internet is not a very attractive prospect. A recent incident on the Internet with the Sun *rwall*d facility in which a user, unintentionally, succeeded in broadcasting a message to nearly every host in the Internet, provoked a storm of protest, for example.

As a result, techniques for limiting broadcasts through the use of multicasting methods, augmented by multicast servers where natural broadcast transmission isn't available, have been developed and are being explored. In RFC 966 and RFC 988, one can find the seeds of an Internet service which seeks to provide the benefits of multicasting to an Internet environment without paying a broadcast price. The recent RFCs on NetBIOS (1001,1002) make reference to Host Group Multicast as a possible extension of the present proposals.

Where is it all going? Just as magazines, newspapers, television, radio and other mass media are useful, in part, because they permit messages to be delivered to unknown parties, the computer communications applications domain will evolve requirements and services based on this concept. The opportunity for creativity (and, perhaps, profit!) lies in discovering efficient, low-cost methods for achieving comparable results.

* [John Shoch, "Internetwork Naming, Addressing and Routing", COMPCON Fall 78 Proceedings, IEEE Catalog Number 78CH-1388-8C, Washington, DC, September 1978].

TCP/IP Interoperability Conference analysis

In March 700 people ventured to Monterey to attend the widely acclaimed TCP/IP Interoperability Conference. This large and diverse collection of people educated each other about the meaning and directions of the TCP/IP suite of protocols. The distribution of attendees was:

2%	Press
4%	Research
12%	Military User
12%	Systems Integrator
14%	Commercial User
18%	University User
38%	Vendor

400 take tutorials

Four hundred people attended the tutorials. Topics covered were TCP/IP for Unix, VMS, VM, and MVS, plus Internet Systems Planning and Engineering (or "The Joy of Gateways") indicating a very high interest in learning about how the protocols work in specific operating environments.

300 organizations

While some organizations were represented by upwards of a dozen people most groups sent one or two persons in order to become more informed about the TCP/IP protocols and services offered by them.

NetBIOS RFCs issued

After nearly a year of effort the vendor community issued a pair of RFCs that specify the method for implementing NetBIOS services on top of TCP/IP. (See page 15 for details.)

Network management next

Network management is becoming extremely important. Today there are a number of products available that do an incomplete job of managing a simple LAN. A customer who purchases different products finds himself in the position of having to also purchase a number of different network management packages. Not only does this have a negative cost impact on the customer but the mere inclusion of (perhaps) overlapping network management packages does not ensure the correct monitoring or operation of the customer's network. A number of vendors have decided to band together to produce written specifications (much like the NetBIOS effort) that will allow any vendor to create a network management product to handle all the workstations, PCs, hosts, bridges, routers, gateways, etc. that live in a customer's environment. The first public meeting is being held in May.

BOFs popular

Perhaps the most exciting meetings were held in eight Birds of a Feather sessions that centered on topics of interest to the various constituencies. BOFs are extremely valuable at a widely attended conference because "all the right players" are available for discussion and resolution. The eight BOFs were:

- Network File Models
- The Politics of Campus/Corporate Networking
- V.3 TCP/IP and STREAMS Design Decisions
- Real-time Network Applications and Protocols
- Gateway Monitoring
- NSFnet + Regional Rumors and Gripes
- Telnet Extensions
- Telnet Option Negotiations to IBM Hosts

Have you heard the GOSIP?

The draft version of GOSIP, Government Open Systems Interconnection Procurement Specification for Fiscal Years 1987 and 1988, was released for comment just before Christmas last year.

What is the scope of GOSIP?

GOSIP, which goes hand-in-hand with the NBS OSI Implementors Agreements, sparked off some considerable discussion, in particular on the Arpanet TCP-IP mailing list. The discussion centered around several topics. To begin with, there was considerable argument about the scope of GOSIP. Most of the participants agree that the applicability section of the document is open to some interpretation. There seems to be a difference between what the document says versus what the authors *intended* it to say. Here is the crucial wording:

"GOSIP is to be used by all Federal Government agencies when procuring ADP systems or services and communication systems or services. It is mandatory for all new network implementations and should be carefully considered for retrofits. Specifically, this specification is mandatory and applies to the procurement of all mini computers and mainframes that are to be interconnected as end systems or intermediate systems. Optionally, it applies to, and is strongly recommended for, all microprocessors, intelligent terminals, work stations, and personal computers that will be interconnected as end systems or intermediate systems.

Although GOSIP mandates OSI implementations in products, it does not preclude the acquisition of additional (perhaps vendor-specific) networking capabilities in that same equipment.

For a period of two years, agencies are permitted to procure alternative interoperable protocols, but they must provide a mechanism for those protocols to interoperate with OSI protocols as well."

Some readers interpret this to mean that anyone wishing to buy networking products now, *must* "go OSI", while others see it as a guideline for agencies who *choose* to select products from the ISO suite. It is expected that the wording will be further clarified in the final version of the document. (Comments were due back to NBS on March 2, 1987.)

Two points of view

In the wake of the "applicability" discussion followed general opinions about the ISO/OSI philosophy. We obtained permission to reprint two messages, which we feel reflect the two camps on this issue. The first, from Phil Karn, focuses on the political forces behind international standards development. The second, by Karl Auerbach, is a defense of GOSIP. In the future we will bring you updates on the GOSIP situation and explore what impact it will have on the adaptation of the ISO protocol suite in the United States.

continued on next page

Have you heard the GOSIP? *(continued)*

CON: From: Phil R. Karn

Subject: GOSIP vs TCP/IP

Indeed, one wonders if the computer public is at all aware of the fact that the Internet has been quietly doing what the ISO hawkers are only promising in big bold trade rag headlines. On the other hand, most vendors have certainly heard of TCP/IP, considering that most of them already sell it, so they have less of an excuse.

Darker forces I think there are other, darker forces at play here. Recent developments (or, more precisely, the lack of same) in the high definition TV standards game illustrate what I think is going on in our own field. It seems that over the past few years, certain Japanese companies have led the way by developing a complete line of compatible, working, high quality video components. You can buy their stuff off the shelf right now. At a recent international standards convention in Europe, the Americans and the Canadians enthusiastically supported the Japanese standard. After all, it works and it's available now. "Can't have that," the Europeans replied. "It'd be too hard to scan-convert back to our existing 625-line 50-Hz formats." And everybody went home without an international standard.

**Good for users or
good for business?** Really now. And they said it with a straight face. This has to be about the thinnest technical excuse anyone has ever invented. The *real* reason (and this was openly stated in a *European* trade journal editorial) is that the European manufacturers deeply resent the Japanese head start into the high definition TV business. There is just no way they are going to approve anybody else's standard, regardless of how good it is technically or whether it's good for the users or not. It'd be bad for business.

To the European vendors, I am truly sorry that the Americans got a head start by inventing TCP/IP and being the first to build big, operational internetworks in which the common carriers ("PTTs") are only minor subcomponents. To the American vendors of protocol software, I am truly sorry that so many public domain implementations of TCP/IP are out there stealing your sales.

To those well-meaning souls in the federal government and elsewhere who naively trust vendor groups and standards organizations to know what's best for your networking needs: take a look at the prices they're charging for the few ISO packages out there. After you've put your eyeballs back into your head, kick out the salesman and take a good close look at just what these slickly advertised protocols will do for you (as distinguished from your vendor's stock price). Then decide if you want to throw everything away and start over just so you can use the magic phrase "ISO compatible" to describe your network.

**TCP/IP was
designed by users** TCP/IP is uniquely successful among communications standards because it was one of the very few ever designed by the *Users* (who just want to get their work done as efficiently and as cheaply as possible) instead of the *Vendors* (who want to make as much money as possible, an entirely different goal). What's good for General Motors may sometimes be good for the country, but in the protocol standards game it's a different story. Michael Padlipsky was right on target on this one. Only the most hopelessly naive user succumbs to the "Illusion of Vendor Support." [Ed: See RFC 873]

**PRO: From: Karl Auerbach ,
Epilogue Technology Corp.**

Subject: A defense of GOSIP

I think GOSIP is a good idea. I support it. I have read GOSIP. I have read the NBS agreements and have participated in the NBS Implementors' Workshops. I have read, and believe I understand, many, if not most of the ISO and CCITT specifications. I have been implementing one of the larger parts (X.400).

1. As for GOSIP mandating a universal government-wide requirement: No matter how one reads the express language of the document, does anyone really think that agencies will abandon SNA? If SNA is an implicit exception why not TCP, XNS, ... ?

**TCP/IP is
changing too**

2. The entire validity of the ISO protocol suite has been called into question because some of the standards have changed as they progressed through the standardization process. So what? Couldn't the same reasoning be applied to the TCP suite because new RFC's are issued?

**No gateway
protocols yet**

Yes, the ISO protocols and services are changing. Our own X.400 implementation will be, in part, invalidated due to changes which will be adopted next year. And is ISO missing important parts? Yes. For instance its protocols for handling routing between intermediate systems ("gateways" in TCP terms) are still being developed. But can one really say that the Internet has done a really good job of inter-gateway routing?

Does MAP/TOP contain some really incredibly dumb ideas? Yes. For example, network level addresses (NSAPs) contain the *physical* media addresses (e.g. the 48 bit Ethernet address). This can become a management nightmare, especially as the NSAP is a necessary component of higher level addresses and will be stored by the various application-level directory services. But this oddity is *not* a necessary part of ISO, only a temporary expedient reasonably adopted by the MAP folks to defer inventing ARP and routing protocols. GOSIP places a high priority on resolving this issue. And answers are presently being considered, just read for example, RFC 995. ["End System to Intermediate System Routing Exchange Protocol for use in conjunction with ISO 8473" - Ed.]

**Grow up with
new technology**

Does this mean that one should not "go ISO"? Perhaps, if you are measuring costs over a short term. But, if you take a long view, and believe that ISO will, in fact, mature, then perhaps you ought to invest now, grow up with the technology, and avoid a conversion expense.

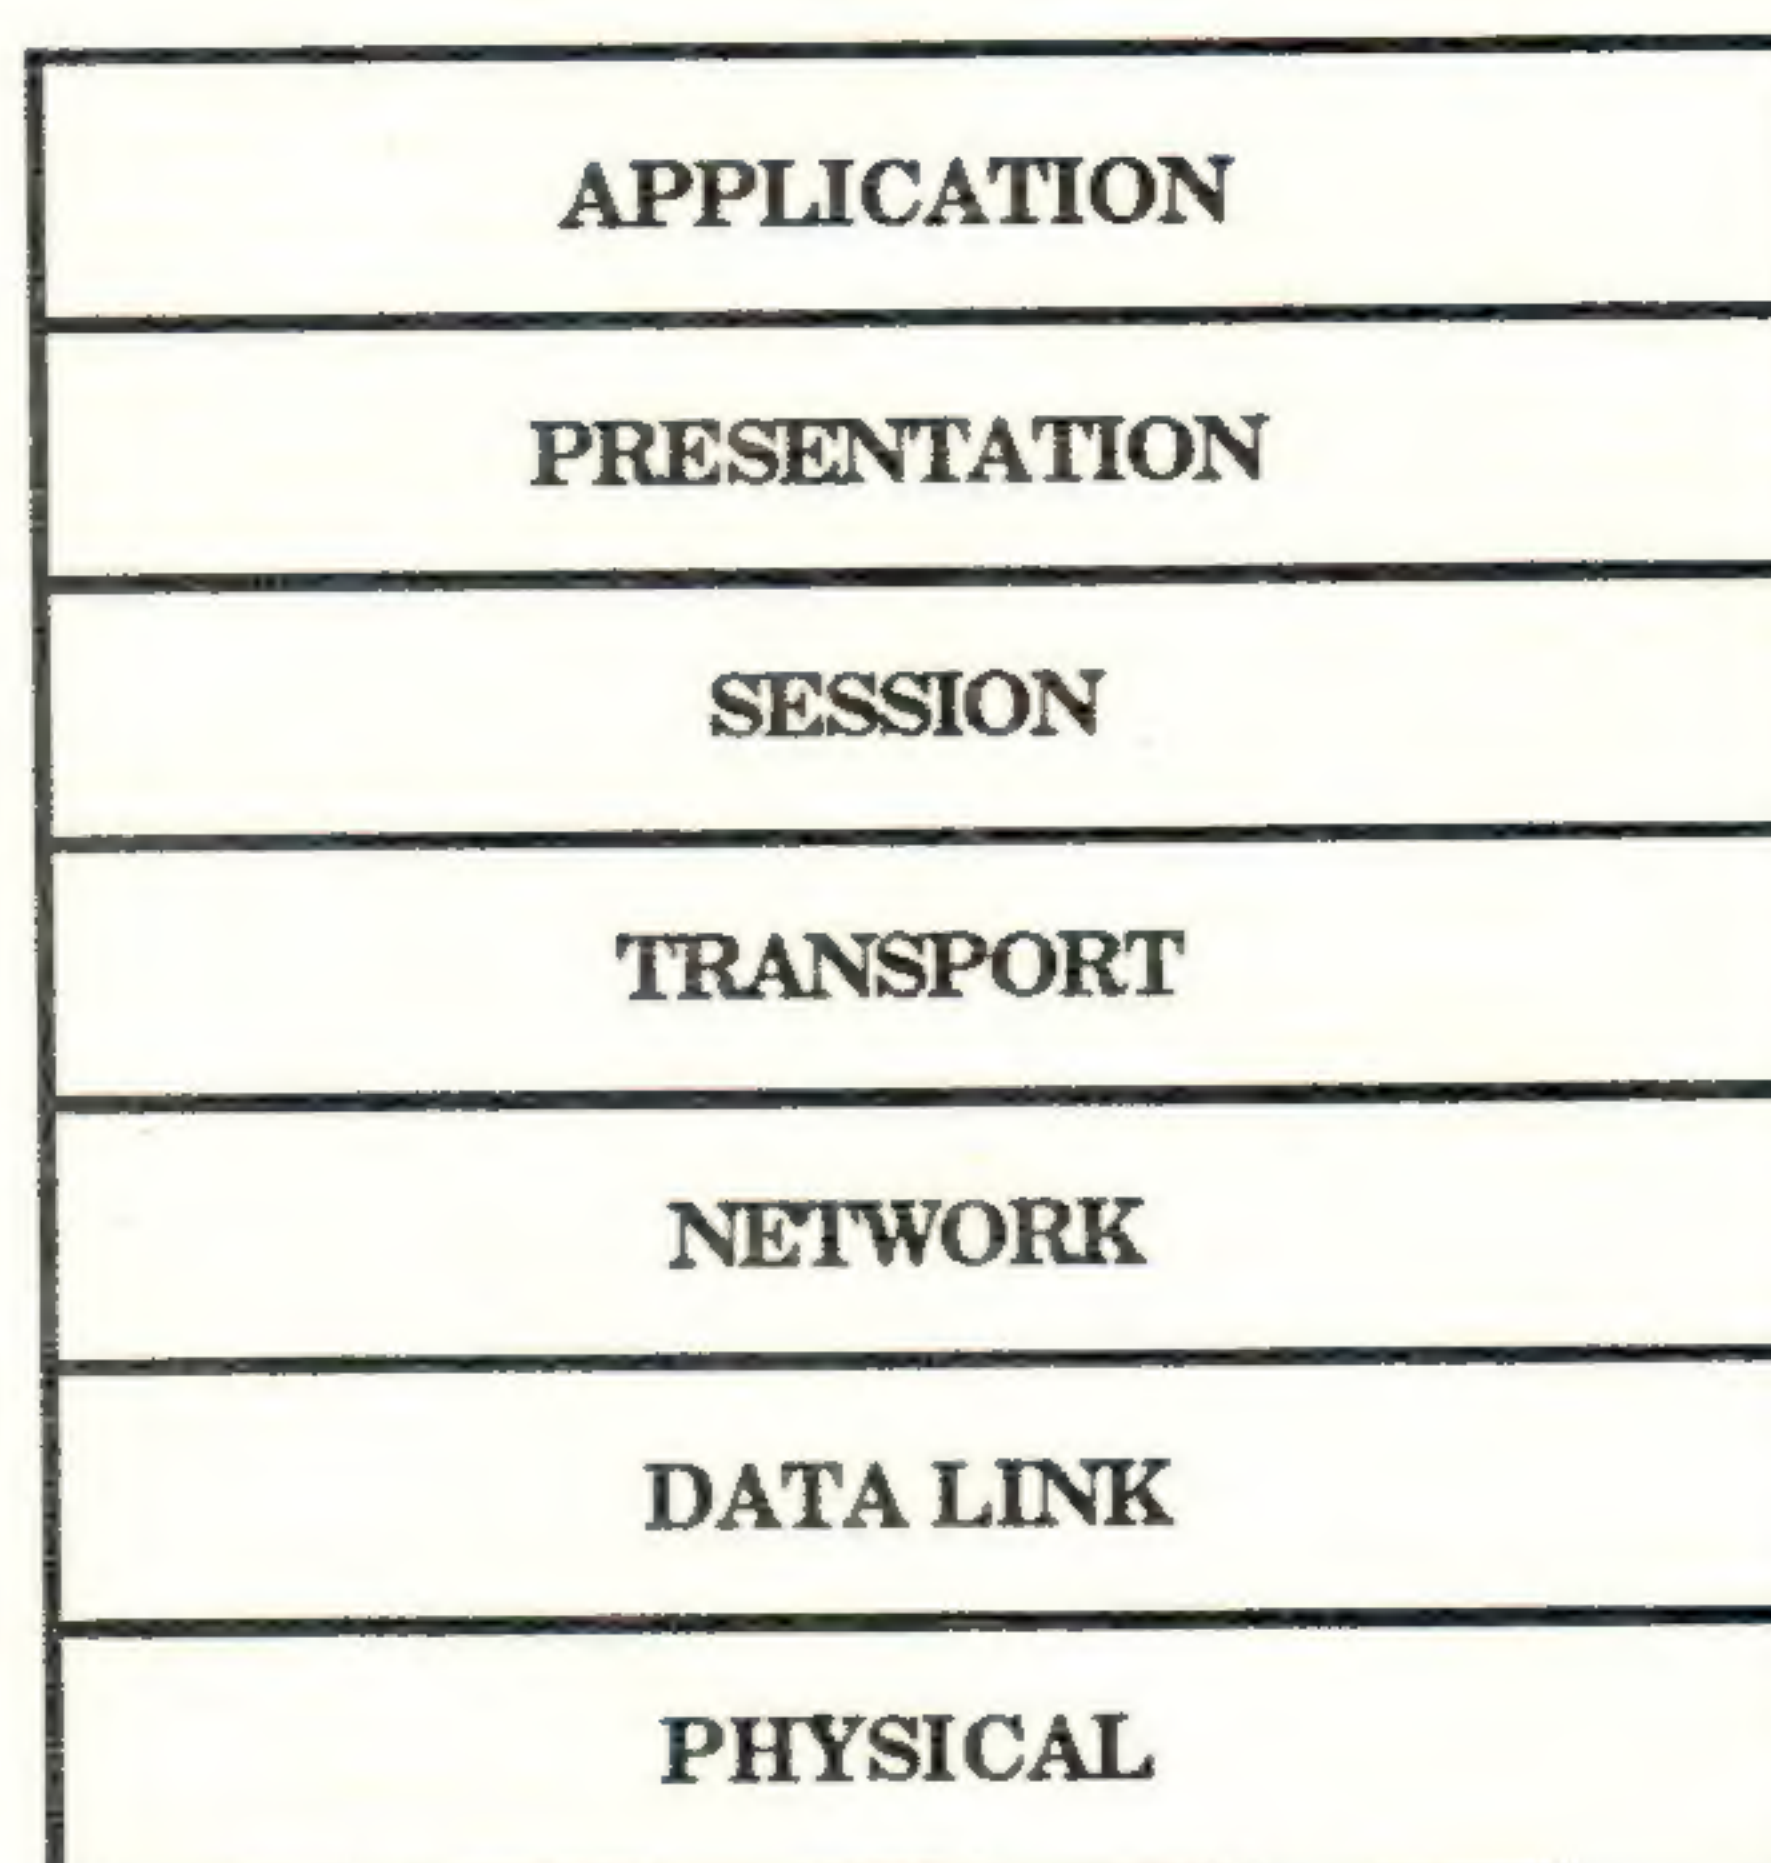
**Implementors'
Agreements limits
the options**

3. The ISO protocols and services contain many, many good ideas. They are in many respects superior to TCP services. There has been criticism that the ISO work is bloated. I believe this is a valid objection. But if you look at the Implementors' Agreements you will find many portions of the full ISO specifications have been chopped off or limited. In addition, as I have worked with ISO my viewpoint has changed. For example, at first I considered most of the session synchronization functions to be questionable. Why should I pay their cost when I am never going to use them? It turns out that they are extremely useful. And, in practice, they don't seem to cost much. I remember similar arguments being raised by assembler language programmers against "the terrible waste of high level languages."

ISODE: Horizontal Integration in Networking

by Marshall Rose, Northrop Research
and Technology Center

The ISO Development Environment at the Northrop Research and Technology Center (ISODE, pronounced *eye-so-dee-ee*) is a system that implements the upper layers of the ISO protocol stack on top of the TCP. In a sense, it combines the best of both worlds. In this article, we consider the rationale for such an undertaking.



The ISO Reference Model

The three upper layers of the ISO architecture which reside above the transport layer are the session, presentation, and application layers.

The *Session* layer, which resides directly above the transport layer, is responsible for managing distinct dialogues on the end-to-end connection. One example of this might be the use of checkpointing during a large data transfer.

Directly above the session layer is the *Presentation* layer, which is responsible for the mapping of machine-dependent data structures to machine-independent representations that are sent over the network. For example, the presentation layer enforces a consistent representation of floating-point numbers across the network.

Service elements

Directly above the presentation layer is the *Application* layer which is really composed of several "service elements". There are typically three present in each application: the association control service element, which opens and closes the connection to the remote system; the remote operations service element, which is a superset of many conventional remote procedure call facilities; and an application-specific element, which directs the protocol between the two application entities.

Abstract Syntax Notation	Integral to the presentation and application layers is the notion of an <i>Abstract Syntax Notation</i> (ASN). This fancy name belies a simple idea: when discussing the data structures transmitted by an application, these data structures are described in a form which can be manipulated by a computer program. The abstract syntax permeates the upper two layers: the presentation layer maps the syntax describing the data structure into a transfer notation which describes how to encode those data structures as bits and bytes. Similarly, in the application layer, ISODE has a compiler, called <i>pepy</i> , which in addition to performing some consistency checking on a specification, also generates a decoder which will read the information from the presentation layer into the "pretty-printers" of this information to aid in debugging. These facilities have proved invaluable in development work thus far.
TCP/IP solves the interoperability problem	Although many herald the ISO protocol suite as ushering in the age of open systems interconnection, the simple fact is that TCP/IP has already accomplished this feat in the marketplace. ISO probably won't be as successful for another 3 - 5 years. The key strength of TCP/IP is in its ability to connect different types of networks. In contrast, as soon as multiple ISO networks are connected together a problem arises because currently there are no routing or gateway-to-gateway protocols in the ISO suite. Further, considering the large number of mature TCP/IP implementations available today, it becomes obvious that TCP/IP has solved the interoperability problem. Of course, the ISO protocol suite will catch up and inevitably surpass TCP/IP. However, the power-curve involved is enormous, and the TCP/IP world continues to gain experience.
Large experimental basis	The DARPA/NSF Internet is a fantastic environment in which to perform internetting experiments. The researchers in this community are constantly learning both from their successes and their failures. In contrast, until the ISO community can build an internet they cannot even <i>conduct</i> experiments!
No general data representation	Although TCP/IP excels in internetting, the applications, although mature, aren't that well developed. The TCP provides a clean end-to-end circuit between two applications, which the applications designer is expected to use "raw". If any type of dialogue facilities are required, then each application uses its own. Similarly, each application has its own application-specific representation of how data is encoded. These practices are both wasteful and counter-productive as they don't foster the use of general facilities which could be used in a wide range of applications. Instead, each application designer is encouraged to "re-invent the wheel," omitting features which might be useful, and develop a new set of tools to debug, test, and measure each new application.

continued on next page

ISODE (*continued*)**Encouraging
performance
results**

One argument in favor of the way TCP/IP applications are currently constructed is that using session, presentation, and application services will result in inefficient systems. The implementation of ISODE shows that these mechanisms can be implemented efficiently. (An application using remote operations tends to run only with 15% less throughput than a corresponding benchmark using raw TCP.) To optimize efficiency, the key concept is to spend the least amount of time moving the most amount of data. It's something that has to be an integral part of the system design, but it's also fairly easy to do in practice. Of course, if a particular facility isn't needed, such as checkpointing, it can be negotiated away.

**UCL is
implementing
X.400 on top of
ISODE**

Experience with ISODE has shown that the upper-layer ISO mechanisms and the applications which use these mechanisms are quite powerful. Further, the experiences of the Internet community can have a substantive impact on the design and implementation of "next generation" systems. Briefly, let's consider one project which illustrates this. A group at the University College London is working on message handling (X.400) and directory service systems. Their applications reside on top of ISODE. The implementors have a lot of experience with computer mail in the DARPA/NSF Internet. In the design of their transport system for multi-media messages, these experiences, along with all of the successes and failures the Internet community has seen in the last 15 years, contribute toward making the final system something spectacular.

**Vertical versus
horizontal
integration**

Next, we explore the controversy regarding the ISODE approach. The argument against an ISODE approach is valid depending on whether networking is viewed as an exercise in *horizontal* or *vertical* integration. (This analysis is due to Einar Stefferud of Network Management Associates.) A protocol "purist" takes the vertical view: the protocol stack, in its entirety, is important. For example, efforts like MAP/TOP (the OSI user's group under SME sponsorship) foster the vertical view. The MAP/TOP user's profile emphasizes a particular slice through the ISO protocol suite. Although the profile chooses particular options at each layer, each layer is chosen as an integral part of the entire stack. This results in a product specification that can be procured from a single source.

In contrast, a protocol "pragmatist" takes the horizontal view: each protocol layer, viewed in terms of the services it provides and requires, is important. If each layer provides the service as defined by its specification, then regardless of its implementational details, the resulting multi-layer system is fully functional. However, care must be taken to make sure that two adjacent layers interface correctly (given consistent service definitions, this is not a problem in practice.)

When implementing ISODE, our intent was to implement as many options as possible in order to provide the widest base for experimentation. This has proved largely successful: far more effort was spent implementing the kernel of each layer than in adding the "bells and whistles" which compose the various options.

Mix and match layers

Herein lies the controversy: at present, since ISODE uses TCP instead of a "real" ISO transport protocol, a current ISODE implementation cannot directly talk to a native ISO implementation. Naturally, one could take the "pure" ISO portion of ISODE (the vast majority of the code) and install it above a native ISO transport implementation. But, this implies multiple vendor sources for various layers. Perhaps though, this is not a problem: one of the claimed advantages of open systems interconnection is to foster a multi-vendor environment; one can argue that this should extend to the protocol layers and not just the protocol stack. In other words, what is the point of all this architecture if layers implemented by different vendors cannot be mixed?

There are good reasons for wanting to use implementations from multiple sources: different vendors tend to excel in different areas of expertise, which encourages a multi-vendor solution when building an optimal multi-layer system. For example: the experience, disposition, and culture required to build an optimal implementation of the ISO transport service is vastly different than what is required to build an optimal implementation of the ISO remote operations service. It may stress the expertise resources of even large vendor organizations to provide both expertly.



Eventually, ISODE will contain a complete ISO protocol stack, so that it can operate as a native ISO environment. ISODE is available for a nominal fee with a trivial license agreement. Once the terms and conditions are agreed upon, the recipient is free to use ISODE in any fashion. This includes selling systems based on ISODE --- even selling such systems back to the Northrop Corporation. This is fine, since presumably any vendor doing this is selling support. Considering that Northrop has sponsored ISODE entirely out of internal R&D funds, this is particularly generous. (Information on how to obtain the ISODE distribution can be found on the next page.)

Only openly available OSI implementation

These arguments explain why ISODE, or something like it, is exactly what is needed in order to leverage the experience of the DARPA/NSF Internet with the potential of the upper-layer ISO architecture. The message handling system described earlier is an excellent example of this. ISODE has been called "the only openly available implementation of open systems interconnection," and now the reasons should be clear.

Marshall T. Rose received the B.S., M.S., and Ph.D. degrees in Information and Computer Science from the University of California, Irvine, in 1981, 1983, and 1984 respectively. He has worked on formal methods of specification and verification of computer network protocols, models of parallel and distributed computation, and computer-based message systems. He currently works in the Computer Science Laboratory at Northrop Research and Technology Center, and is also an Adjunct Assistant Professor at the University of Delaware.

About the ISODE 2.0 software

Non-proprietary ISODE is not proprietary, but it is not in the public domain. This was necessary to include a "hold harmless" clause in the release. The upshot of all this is that anyone can get a copy of the release and do anything they want with it, but no one takes any responsibility whatsoever for any (mis)use.

UNIX now, others later ISODE runs on native 4.2BSD and SVR2 with an Excelan card. It also runs on a AT&T 3B2 running SVR2 and the WIN TCP/IP package, native ROS (the Ridge Operating System), and HP-UX. (Future releases will support VAX/VMS, PC-XENIX, and a variant of PC/IP.)

The discussion group ISODE@NRTC.NORTHROP.COM is used as an open forum on ISODE.

Much documentation The primary documentation for this release consists of a User's Manual (approx 300 pages) and a set of UNIX manual pages. The sources to the User's Manual are in LaTeX format. In addition, there are a number of notes, papers, and presentations included in the documentation set, again in either LaTeX or SLiTeX format.

For more information, contact:

Northrop Research and Technology Center
Attn: Automation Sciences Laboratory (0330/T30)
One Research Park
Palos Verdes Peninsula, CA 90274
USA +1-213-544-5393

Getting ISODE There are several ways to get a distribution:

DARPA/NSF Internet: If you can FTP to the DARPA/NSF Internet, you can use anonymous FTP to louie.udel.edu [10.0.0.96] and retrieve the file portal/isode-2.tar.

NIFTP: If you run NIFTP over the public X.25 or over JANET, and are registered in the NRS at Salford, you can use NIFTP with username *guest*, and your own name as password, to access UK.AC.UCL.CS to retrieve the file <SRC>isode-2.tar.

Mailings:

In North America, send a check or invoice for \$130 US to:

Department of Electrical Engineering
Attn: Prof. David J. Farber
University of Delaware
Newark, DE 19716
USA +1-302-451-1163

In Europe, send a cheque or invoice for £100 sterling to:

Department of Computer Science
Attn: Soren Sorenson
University College
Gower Street
London, WC1E 6BT
UK +44-1-387-7050 x3680

The *gated* program - A user's review

by Sergio Heker, John von Neumann
National Supercomputer Center

The *gated* program which was developed by Mark Fedor of the Cornell Theory Center is currently running on John von Neumann Center's network (JVNCnet) and its associated gateways to the members of the JVNC Consortium.*

Replaces *routed* and *egpup*

gated provides the ability to interchange routing information using the HELLO, RIP and EGP protocols amongst the gateways. *gated* was designed as a replacement for the existing UNIX 4.2BSD *routed* and *egpup* as well as the implementation under UNIX of the HELLO protocol. It allows for maximum configurability and interaction between the gateway protocols without violating the protocol specifications. The design of the program allows the user to specify what kind of routing information and protocol type to listen for on a per-interface basis. This ability to specify interface/protocol is especially valuable in an environment like JVNC where we have to talk HELLO with the NSFnet peers, RIP with our gateways, and EGP with some of our routers.

Avoiding routing loops

JVNC's major routing problem has been the lack of a mechanism to avoid routing loops. JVNCnet connects 16 universities; many of them are on the Arpanet and some are connected via other nets to other NSFnet sites. If not handled carefully, false routing information leading to loops could potentially take down the entire NSFnet. We are extra careful in what networks we announce to the world and what kind of routing metrics we use. As an NSFnet backbone site we are responsible for the announcement of networks to the backbone. *gated* handles these announcements in two different ways. In the first form, all networks which are to be announced (together with interface, protocol and routing metric information) are entered into a configuration file. Under this scheme, no other networks will be announced to the NSFnet backbone. The second form allows announcement of all networks *except* the ones listed in the configuration file. At JVNC we use the first form since this gives us tighter control of our complex environment. In order to maintain consistent information it is necessary for the various NSFnet sites that run *gated* to cooperate. An effort is under way to create a centralized file containing information about "allowed" networks. JVNC currently contributes approximately 20 networks to this central configuration file.

* The members of the JVNC Consortium are: MIT, Harvard, Penn State, University of Pennsylvania, Princeton, Institute for Advanced Studies, Rutgers, Columbia, NYU, University of Rochester, University of Colorado, University of Arizona, Stevens Institute of Technology, New Jersey Institute of Technology, University Medical and Dental Center of New Jersey.

continued on next page

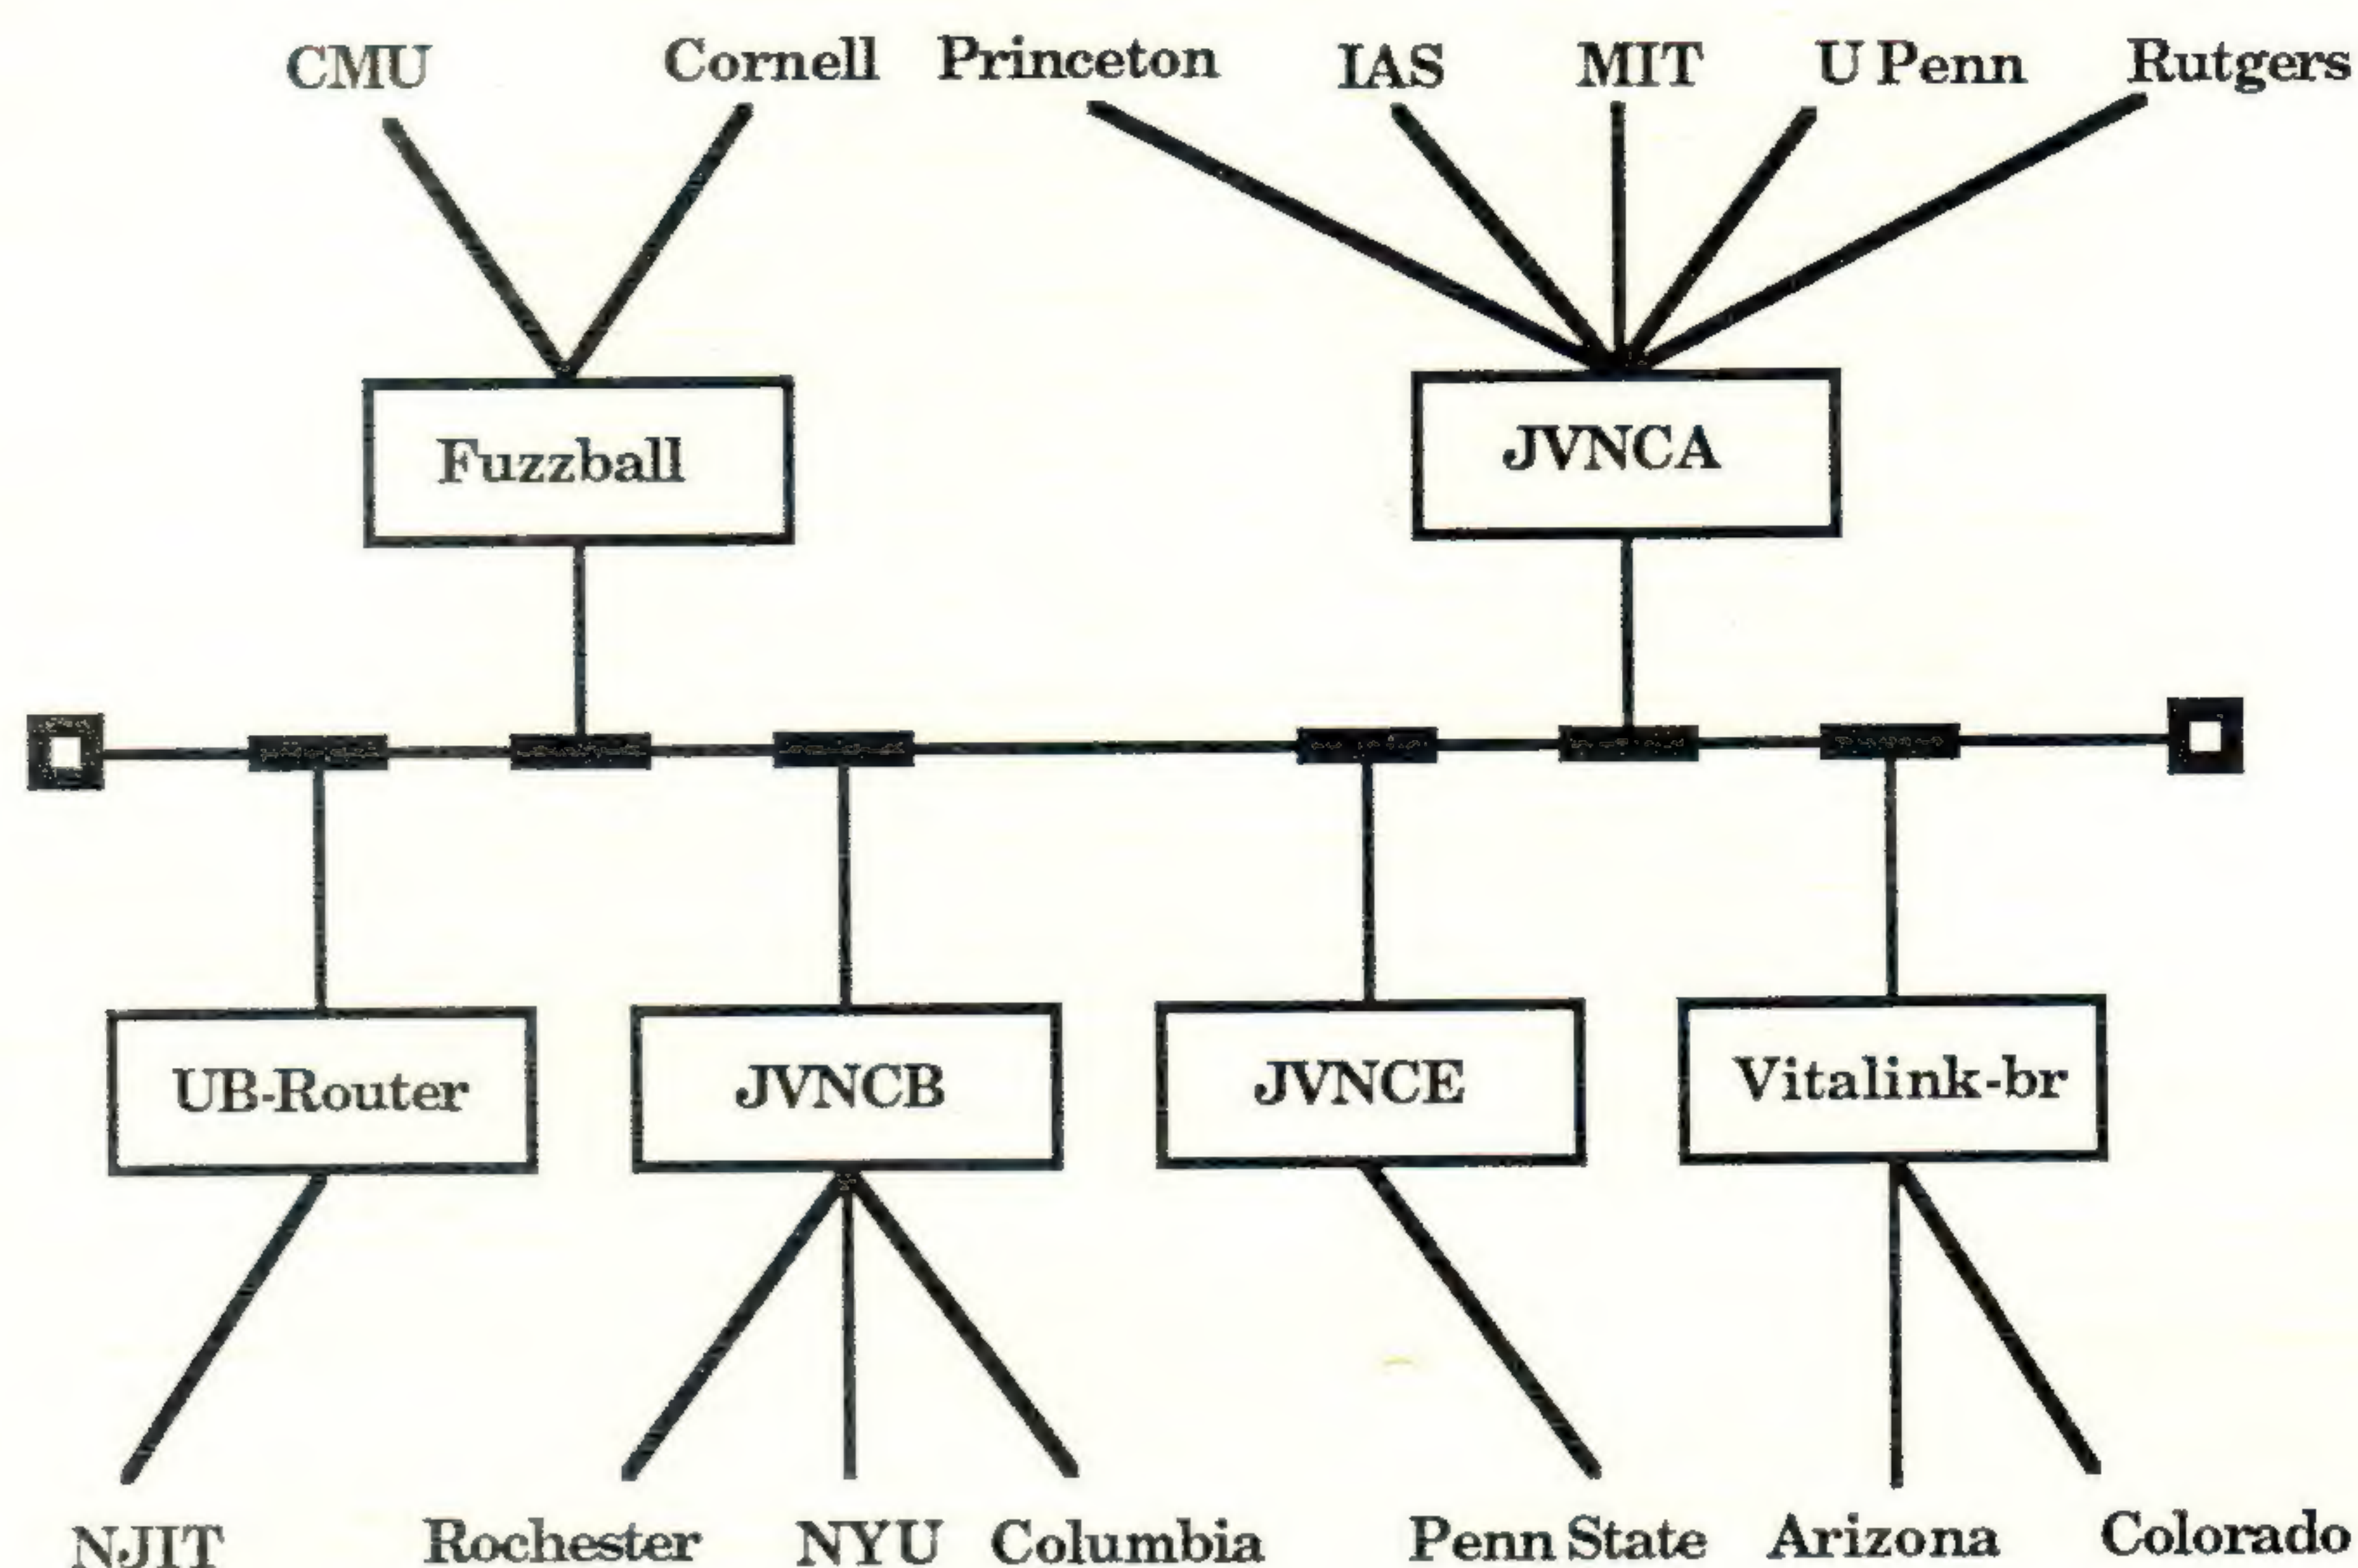
gated - A user's review (continued)

Figure 1.

Flexible re-routing

Part of the JVNCnet is shown in Figure 1. All the sites on the local area network speak RIP with one another and with the remote gateway, EGP with the UB-router, and HELLO with the Fuzzball. Several of the sites shown bridge other networks which makes alternate paths possible. We are able to observe the change in routes and can make configuration changes to accommodate for different network conditions. For instance, when JVNCB's high speed serial link to Columbia University is down, all traffic between JVNC and Columbia is rerouted either through the Arpanet or via NYSERnet/NSFnet/JVNC. This is an example of the flexibility of *gated*.

The current status of the dynamic routing implementation in our network is that most of the JVNC sites are running *gated*, and a few are in the process of installing *gated* and integrating local campuses to JVNCnet and the other regional and national networks. We cooperate with the members of our consortium in the distribution, installation, and updating of *gated* when we receive the code from Cornell. Our experience with *gated* has been a good one. At the same time, we do not see any other way of managing the routing configuration in our complex environment.

We are very thankful to the Cornell Theory Center and in particular to Mark Fedor for their support and the opportunity of being one of the first networks to test the *gated* software.

For more information on JVNCnet, send electronic mail to Sergio Heker: heker@jvncc.csc.org.

NetBIOS RFCs issued

At the March TCP/IP Conference in Monterey, Excelan, Sytek, and Ungermann-Bass presented two new RFCs describing the specifications for NetBIOS on top of TCP/IP. It was the first time in history that *vendors* got together to write RFCs, and as such it marks the beginning of a new era for the TCP/IP protocol suite. The two documents are:

RFC 1001: Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods

RFC 1002: Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications.

As always, these documents can be copied via FTP from the RFC: directory on the SRI-NIC.ARPA host. (RFC:RFC100x.TXT). You may also order hardcopies for \$5.00 each from:

DDN Network Information Center
SRI International, Room EJ291
333 Ravenswood Avenue
Menlo Park, CA 94025
800-235-3155 or 415-859-3695

Call for Participation

The 1987 ACM SIGCOMM Workshop on Frontiers in Computer Communications Technology, August 11-13 in Stowe, Vermont is an international forum for those interested in the theory, development and applications of computer communications. At the workshop there will be a session on TCP/IP interoperability. Topics that will be covered include: TCP instead of TP4, efficiency of TCP implementations, TCP/IP in Europe and Japan. This session is being chaired by Dan Lynch of Advanced Computing Environments. If you wish to participate as either a speaker or attendee please contact him at 408-996-2042. For more information about the workshop contact:

J.J. Garcia-Luna, Program Chair
SRI International, Room EK319
333 Ravenswood Avenue
Menlo Park, CA 94025
415-859-5647
Internet: garcia@istc.sri.com
Telex: 334486

UNIX is a trademark of AT&T Bell Laboratories.

MS is a trademark of Microsoft Corporation.

VMS is a trademark of Digital Equipment Corporation.

IBM is a trademark of International Business Machines Corporation.

Xenix is a trademark of Microsoft Corporation.

CONNE~~X~~IONS

21370 Vai Avenue
Cupertino, CA 95014

CONNE~~X~~IONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Associate Director, University of Southern California Information Sciences Institute.

Subscribe to CONNE~~X~~IONS

U.S./Canada \$360. for 12 issues/year
University \$240. for 12 issues/year
International \$ 50. additional per year

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNE~~X~~IONS). ☐ Bill me/PO# _____

☐ Charge my ☐ Visa ☐ Master Card Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNE~~X~~IONS

21370 Vai Avenue
Cupertino, CA 95014

CONNE~~X~~IONS